

07-17-00 A

**UTILITY
PATENT APPLICATION
TRANSMITTAL**
(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))

Attorney Docket No.	112025-0067C1
First Inventor or Application Identifier	Arthur Lin et al.
Title	ROUTER WITH CLASS OF SERVICE MAPPING
Express Mail Label No.	EL056746535US

APPLICATION ELEMENTS
See MPEP chapter 600 concerning utility application contents

1. *Fee Transmittal Form (e.g., PTO/SB/17)
(Submit an original and a duplicate for fee processing)

2. Specification [Total Pages **17**]
(preferred arrangement set forth below)

- Descriptive title of the Invention
- Cross References to Related Applications
- Statement Regarding Fed sponsored R & D
- Reference to Microfiche Appendix
- Background of the Invention
- Brief Summary of the Invention
- Brief Description of the Drawings (if filed)
- Detailed Description
- Claim(s)
- Abstract of the Disclosure

3. Drawing(s) [Total Sheets **5**]
4. Oath or Declaration [Total Pages]

a. Newly executed (original copy)
b. Copy from a prior application (37 C.F.R. § 1.63(d))
(for continuation/divisional with Box 17 completed)

[Note Box 5 below]

DELETION OF INVENTOR(S)
Signed statement attached deleting
inventor(s) named in the prior application,
see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).

5. Incorporation By Reference (useable if Box 4b is checked)
 The entire disclosure of the prior application, from which a
copy of the oath or declaration is supplied under Box 4b, is
considered to be part of the disclosure of the
accompanying application and is hereby incorporated by
reference therein

Assistant Commissioner for Patents
ADDRESS TO:
Box Patent Application
Washington, DC 20231

6. Microfiche Computer Program (Appendix)
7. Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)

- a. Computer Readable Copy
- b. Paper Copy (Identical to computer copy)
- c. Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. Assignment Papers (cover sheet & document(s))
9. 37 C.F.R. §3.73(b)
9. Statement (when there is an assignee) Power of Attorney
10. English Translation Document (if applicable)
11. Information Disclosure Statement (IDS)/PTO-1449 Copies of IDS Citations
12. Preliminary Amendment
13. Return Receipt Postcard (MPEP 503)
(Should be specifically itemized)
*Small Entity
14. Statement(s) Statement filed in prior application,
((PTO/SB/09-12)) Status still proper and desired
15. Certified Copy of Priority Document(s)
(if foreign priority is claimed)
16. Other:

"NOTE FOR ITEMS 1 & 14: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. §1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28)."

17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: 09/900,478

Prior application Information: Examiner

Group/Art Unit:

18. CORRESPONDENCE ADDRESS

<input type="checkbox"/> Customer Number or Bar Code Label	(Insert Customer No. or Attach bar code label here)			<input checked="" type="checkbox"/> Correspondence address below	
Name	Patricia A. Sheehan				
Address	Cesari and McKenna 30 Rowes Wharf				
City	Boston	State	MA	Zip Code	02110
Country	U. S. A.	Telephone	(617) 951-2500	Fax	(617) 951-3927

Name (Print/Type)	Patricia A. Sheehan	Registration No. (Attorney/Agent)	32,301
Signature	Patricia A. Sheehan		Date
			July 14, 2000

1C008 U.S. PRO
09/616842
07/14/00



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re The Application of:)
Arthur Lin et al.)
)
Serial No.: not yet assigned) Examiner: not yet assigned
)
Filed: herewith) Art Unit: not yet assigned
)
For: ROUTER WITH CLASS OF SERV-)
ICE MAPPING)
)
)

Cesari and McKenna, LLP
30 Rowes Wharf
Boston, MA 02110
July 14, 2000

EXPRESS-MAIL DEPOSIT

"Express Mail" Mailing-Label Number: EL056746535US

The following papers are being deposited with the United States Postal Service
"Express Mail Post Office to Addressee" service pursuant to 37 C.F.R. §1.10:

Continuation Patent Application (17 pages including 29 claims)
 Oath or Declaration (copy) Fee Transmittal Letter
 Utility Patent Application Transmittal Formal Drawing (5 sheets)
Letter
 Check No. 3459 for \$690 Preliminary Amendment

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re The Application of:)
Arthur Lin et al.)
)
Serial No.: not yet assigned) Examiner: not yet assigned
)
Filed: herewith) Art Unit: not yet assigned
)
For: ROUTER WITH CLASS OF)
SERVICE MAPPING)
)
) Cesari and McKenna, LLP
) 30 Rowes Wharf
) Boston, MA 02110
July 14, 2000

Honorable Assistant Commissioner for Patents
Washington, D.C. 20231

Sir:

PRELIMINARY AMENDMENT

Please cancel claims 1-29 and insert the following new claims 30-44 as set forth
below:

30. A router for use in routing packets over a network, the router supporting a plurality, X, of classes of service and including:
 - A. a plurality of input ports for receiving packets over the network;
 - B. a plurality of output ports for transferring packets over the network;
 - C. a classifier for assigning packets received by the input ports to $X * Y$ classes of service, where * represents multiplication, and mapping the XY classes of service to the X classes of service that are supported by the router, the classifier assigning to the

packet one of Y associated levels of priority, wherein each level of priority is associated with a different probability of packet loss;

- D. a buffer subsystem for retaining the packets in class of service per output port queues based on probabilities of discard associated with the X * Y classes of service; and
 - E. a scheduler for transferring the packets from the buffer through each of the output ports based on the X classes of service.

31. The router of claim 30 wherein the buffer subsystem includes multiple storage locations and links available storage locations in a free queue.

32. The router of claim 31 wherein the buffer subsystem includes a processor that determines:

- i. a new weighted average depth for the free queue, and
- ii. a probability of discard for a given packet if the new weighted average queue depth falls below a predetermined maximum threshold associated with the class of service to which the packet is assigned by the classifier.

33. The router of claim 32 wherein the buffer subsystem discards a given packet if the associated new weighted average depth for the free queue falls below a minimum threshold associated with the class of service to which the packet is assigned.

34. The router of claim 33 wherein the buffer subsystem processor calculates the probability of discard as $P_d = c - (m * A_{NEW})$ where c is an intercept and m is a slope that is associated with a line that plots average free queue depth versus probability of discard for the class of service to which the packet is assigned, and A_{NEW} is the new weighted average depth of the free queue.

35. The router of claim 34 wherein the buffer subsystem processor calculates the new weighted average depth of the free queue as $A_{NEW} = A_{CURRENT} + w(I - A_{CURRENT})$ where w

is a weighting factor, I represents the instantaneous depth of the free queue and A_{CURRENT} is the current weighted average depth of the free queue.

36. The router of claim 35 wherein the scheduler selects from the buffer subsystem packets for transfer based on weighting factors associated with the respective 2ⁿ classes of service.

37. A router for use in routing packets over a network, the router supporting a plurality, X, of classes of service and including:

- A. a plurality of input ports for receiving packets over the network;
- B. a plurality of output ports for transferring packets over the network;
- C. a multiple storage location buffer for retaining packets to be transferred through the output ports;
- D. a buffer subsystem for retaining the packets in class of service per output port queues based on probabilities of discard associated with X*Y classes of service, where * represents multiplication; and
- E. a scheduler for transferring the packets from the buffer subsystem through each of the output ports based on the X classes of service that the router supports.

38. The router of claim 37 further including a classifier for:

- i. assigning packets received by the input ports to X*Y classes of service,
- ii. associating the packets with the X*Y classes of service that are supported by the router, and
- iii. assigning to the packet one of X*Y associated levels of priority, wherein each level of priority is associated with a different probability of packet loss.

39. The router of claim 38 wherein the buffer subsystem includes a processor that determines

- i. a new weighted average queue depth for a free queue that links available buffer storage locations, and

ii. a probability of discard for a given packet if the new weighted average free queue depth falls below a predetermined maximum threshold associated with the class of service to which the packet is assigned.

40. The router of claim 39 wherein the buffer subsystem processor calculates the probability of discard as $P_d = c - (m * A_{NEW})$ where c is an intercept and m is a slope that are associated with a line that plots average free queue depth versus probability of discard for the class of service to which the packet is assigned, and A_{NEW} is the new weighted average depth of the free queue.

41. The router of claim 40 wherein the buffer subsystem processor calculates the new depth of the weighted average free queue as $A_{NEW} = A_{CURRENT} + w(I - A_{CURRENT})$ where w is a weighting factor, I represents the instantaneous depth of the free queue and $A_{CURRENT}$ is the current weighted average depth of the free queue.

42. The router of claim 41 wherein the buffer subsystem discards a given packet if the new weighted average free queue depth falls below a minimum threshold associated with the class of service to which the packet is assigned.

43. The router of claim 41 wherein the buffer subsystem retains a given packet if the new weighted average free queue depth is above a maximum threshold associated with the class of service to which the packet is assigned.

44. The router of claim 37 wherein the scheduler selects packets for transfer through each output port based on weighting factors associated with the respective 2^n classes of service.

REMARKS

Please cancel pending claims 1-29 and include claims 30- 44 before examination.

Please charge any additional fee occasioned by this paper to our Deposit Account
No. 03-1237.

Respectfully submitted,



Patricia A. Sheehan
Reg. No. 32,301
CESARI AND MCKENNA, LLP
30 Rowes Wharf
Boston, MA 02110-2699
(617) 951-2500

UNITED STATES PATENT APPLICATION

of

Thomas Edsall

Tauseef Hashmi

and

Arthur Lin

for a

ROUTER WITH CLASS OF SERVICE MAPPING

ROUTER WITH CLASS OF SERVICE MAPPING

FIELD OF INVENTION

The invention relates generally to routers and switches and, more particularly, to routers and switches that support multiple classes of service for packet routing.

5

BACKGROUND OF THE INVENTION

At network multiplexing points, such as switches or routers, the handling of frames or packets is generally determined by rules associated with classes of service to which given frames or packets are assigned. The classes of service essentially define acceptable packet or frame delays and probabilities of packet or frame loss. (As used 10 herein, the term "packet" refers to both frames and packets, and the term "router" refers to both switches and routers.)

The packets are typically assigned to classes of service based on information contained in the packet and/or traffic management rules established by either a network supervisor or a service provider. All packets assigned to the same class receive the same 15 treatment. Being assigned to a "higher" class ensures that a packet will have a shorter maximum transmission delay and a lower probability of loss. Being assigned to a "lower" class may mean a longer delay and/or a greater probability of loss.

Generally, the router maintains at each output port a buffer for holding packets in 20 queues associated with the classes of service. The queues ensure that packets are delivered in order within the various classes of service, and that the associated rules for maximum delays and probabilities of loss can be enforced. Since each queue is essentially separately maintained, the more classes the router supports the more processing and storage capacity is required for a given number of output ports. To

support “x” classes, for example, the router must set aside buffer storage locations for each of the x queues at each of its “y” ports. Further, it must determine for each queue whether or not a next packet should be retained or discarded. The router thus makes $x*y$ separate calculations based on queue length and/or available associated storage locations 5 to determine whether to retain or discard the packets, where “*” represents multiplication.

Network standards, such as the (revised) 802.1p standard, have relatively recently increased the number of classes of service to eight classes. Routers operating under prior standards support four classes of service, and thus, must be upgraded, for example, with increased storage capacities of the output port buffers, to support the increased number of 10 classes. Such upgrading may be prohibitively expensive and/or it may not be feasible. Accordingly, what is needed is a mechanism to operate a router that supports a relatively small number of classes of service in an environment in which packets are assigned to a greater number of classes. Such a mechanism should, without requiring the enlarged 15 storage and processing capabilities conventionally associated with supporting the greater number of classes, maintain service distinctions associated with the greater number of classes and more importantly retain the order of packets within each of the greater number of classes.

SUMMARY OF THE INVENTION

A router maps packets assigned to 2^{n+m} classes of service into 2^n classes of service 20 and assigns the packets to 2^m levels of loss-priority within each of the 2^n classes. The router includes a classifier that uses n bits of an (n+m)-bit “class of service identifier” to map the packets to the 2^n classes, and the remaining m bits to assign the loss priorities. The router then controls packet retention/discard with a modified weighted random early detection scheme based in part on the 2^{n+m} classes and in part on the 2^n classes, to 25 maintain the probability of loss distinctions and in-order packet handling associated with the 2^{n+m} classes.

A scheduler controls the transmission of packets by each output port based on the 2^n classes of service. The scheduler uses a weighted round robin scheme to ensure that

packets from each of the classes are transmitted by each of the output ports within the prescribed maximum delay limits associated with the 2^{n+m} classes of service.

The router includes an output buffer that holds the packets for all of the router's output ports. The router maintains a "free queue," which links the buffer storage locations available for packet storage. To determine whether to retain or discard a given packet, the router compares a weighted average depth of the free queue with predetermined maximum and minimum thresholds that are associated with the particular one of the 2^{n+m} classes of service to which the packet is assigned. If the weighted average exceeds the associated maximum threshold, the router retains the packet in a storage location that is then removed from the free queue and linked to a class of service per output port queue that corresponds to the class of service to which the packet is mapped by the classifier. If the weighted average depth falls below the associated minimum threshold, the router discards the packet. If the weighted average depth falls between the associated minimum and maximum thresholds, the router calculates a probability of 10 discard and compares the probability to a "random" value. The router discards the packet if the probability exceeds the random value, and otherwise retains the packet.

The maximum and minimum thresholds are set relative to one another such that the loss priorities associated with the 2^{n+m} classes are maintained. As discussed below, the router makes only one weighted average queue depth calculation for the free queue, and uses this calculation to determine whether to retain or discard packets for the 2^n classes of service. This is in contrast to prior known routers that must maintain at each output port separate average queue depths for each of the class of service per port queues.

BRIEF DESCRIPTION OF THE DRAWINGS

25 The invention description below refers to the accompanying drawings, of which:
Fig. 1 is a functional block diagram of a network that includes routers that are constructed and operate in accordance with the invention;
Fig. 2 is a functional block diagram of a router of Fig. 1;

Fig. 3 illustrates a mapping of packets to classes of service;
Fig. 4 is a flow chart of the operations of the router of Fig. 2; and
Fig. 5 is a graph of weighted average queue depth versus probability of packet discard.

5 **DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT**

Referring now to Fig. 1, a network 10 includes a plurality of endstations 12 and nodes 14 that transmit packets to other endstations 12 and nodes 14 through routers 16 and 17. The endstations 12 and nodes 14 assign packets to classes of service based on information contained in the packets and/or on predetermined traffic management rules
10 that are provided by the network manager and/or various service providers. The classes of service are essentially associated with maximum limits for transmission delays and probabilities of packet loss. Higher classes are associated with shorter maximum delays and lower probabilities of packet loss. Packets that must be delivered as essential parts of a transmission are, for example, assigned to a higher class than are packets that contain
15 non-essential information.

Preferably, the endstation 12 or the node 14 that introduces the packet to the network assigns the packet to one of 2^{n+m} classes of service. To inform the routers 16 and 17 of the assignment, the endstation 12 or node 14 writes an appropriate class of service “tag” to a COS identifier field in the header that is included in the packet. The COS
20 identifier field has three bits, as defined by (revised) standard 802.1p, and the packet is thus assigned to one of eight classes of service, i.e., 1 of 2^3 classes. The packet is then forwarded by the endstation 12 or node 14 over the network 10 to an input port 28 of a router 16 or 17. The router then transfers the packet through an output port 30 and over the network in accordance with the transmission rules and delay limits associated with the
25 class of service to which the packet is assigned.

The routers 17 support 2^{n+m} classes of service while the routers 16 support 2^n classes, where $n < 3$. We discuss herein the operations of the routers 16 to assign packets to the various classes. Further, we discuss the operations that the routers 16 and/or the

routers 17 perform to determine whether to retain or discard a packet and/or when to transmit a packet.

Referring now to Fig. 2, a router 16 includes a classifier 18 that associates a received packet with one of 2^{n+m} classes of service, based primarily on the COS tag, if any, included in the packet header. The classifier maps the 2^{n+m} classes of service to the 2^n classes based on, for example, the highest order n bits or the lowest order n bits of the COS tag. The classifier then uses the remaining bits of the COS tag to set the loss priorities of the packets. As discussed below, the loss priorities determine if respective packets are discarded or retained during times of congestion. The higher the loss priority of a packet, the less likely the packet will be retained.

If the endstation 12 or node 14 that introduces the packet to the network does not support the 802.1p standard, the COS tag may not be included in the packet. The classifier 18 may then assign the packet to one of the 2^{n+m} classes, currently 2^3 classes, based on appropriate network or service provider transmission rules. It may, for example, assign the packet to a “best effort” class. Alternatively, the router 16 may assign the packet to a particular class of service based on a media access control, or MAC, address included in the packet. The classifier then writes the appropriate COS tag to the packet header.

Referring now also to Fig. 3, the router 16 in this exemplary embodiment supports four classes of service, i.e., 2^2 classes. The classifier 18 maps each of the 2^3 classes of service to an appropriate one of the 2^2 classes of service based on the two highest order bits of the 3-bit COS tag. The third, or lowest order bit, is then used to assign a loss priority to the packet. The classifier 18 thus associates a packet that is assigned to class of service 010 with class of service 01 and sets the loss priority of the packet to 0. Further, the classifier 18 associates a packet that is assigned to class of service 011 with class 01 and sets the loss priority of this packet to 1.

Referring further to Fig. 4 once the classifier 18 associates the packets with the various 2^n classes of service and sets the loss priorities (steps 400-402), a policer 20

enforces network or service provider usage parameter controls by marking, discarding or passing the packets (step 404). The usage parameter controls are set by a network manager or service provider based on, for example, levels of service purchased by or associated with a user. The user may, for example, purchase a level of service based on 5 the transmission of a maximum number of packets per hour. If the number of packets being sent by the user exceeds this limit, the policer then marks, discards or passes the excess packets depending on the traffic management rules.

If the policer 20 marks an offending packet, it assigns the packet to a higher loss priority within the associated class of service. This increases the likelihood that the 10 packet will be discarded if the network becomes congested. In the example, the policer sets the lowest order bit of the COS tag to 1. If the packet is already assigned the highest loss priority within the class of service, the policer 20 either passes or discards the packet, depending on the traffic management rules. If the packet is passed, the policer 20 may charge the user for the use of excess bandwidth.

15 As discussed above, the policer 20 operates in accordance with traffic management rules established by a network manager or service provider. In the embodiment described herein the policer determines if a packet exceeds an established limit by using a “jumping window policing scheme.” The policer thus sets a police rate of B/T for a user, where B is a burst size and T is a time interval and both B and T are set 20 by the network manager or service provider. The policer then counts the number of octets received from the user over intervals of length T. If the count exceeds B, the arriving packet is marked, passed or dropped, depending on the enforcement mode utilized by the policer. Various limits may be set, such as, for example, limits that vary based on the number of times the associated policing rate is exceeded by a given user 25 and/or based on the various classes of service.

A WRED processor 22 determines which of the remaining packets, i.e., the packets that the policer has not discarded, are to be retained in a buffer 24 that holds the

packets for every output port 30 (steps 406-416). The use of a single buffer is in contrast to prior known routers that use a separate buffer for each output port.

The WRED processor 22 utilizes a modified weighted-random early detection (WRED) scheme. The WRED processor associates with each of the 2^{n+m} classes of service, " C_i ," two thresholds, namely, a maximum threshold MAX_{C_i} and a minimum threshold MIN_{C_i} . As discussed below, the thresholds are used by the processor 22 to determine whether to retain or discard a given packet.

The WRED processor 22 keeps track of an average "free queue" depth, which is an average number of available storage locations in the buffer 24. When the buffer is empty, all of the buffer storage locations are linked to the free queue. As packets are retained, buffer locations, which are generally referred to in 512 byte pages, are removed from the free queue and linked to appropriate class of service per output port queues. When the packets are later transmitted, the buffer locations are removed from the class of service per output port queues and again linked to the free queue.

Each time a packet is received, the WRED processor 22 determines a new weighted average free queue depth A_{NEW} :

$$A_{NEW} = A_{CURRENT} + W(I - A_{CURRENT})$$

where I is the instantaneous size of the free queue, W is the weighting factor and $A_{CURRENT}$ is the current weighted average free queue depth (step 406). The weighting factor W is preferably selected such that multiplication is accomplished by shifting the difference value ($I - A_{CURRENT}$). The value $A_{CURRENT}$ is updated at regular intervals with the value of A_{NEW} , such as after every 64B frame time, which approximates the average packet arrival time.

The WRED processor compares the weighted average A_{NEW} with the MAX_{C_i} and MIN_{C_i} values associated with the appropriate one of the 2^{n+m} classes of service. If the weighted average exceeds the MAX_{C_i} value, the WRED processor 22 retains the packet (step 408). If the weighted average falls below the MIN_{C_i} value, the WRED processor 22

discards the packet (step 410). If, however, the average falls between MAX_{c_i} and MIN_{c_i} values, the WRED processor calculates a probability of discard, P_D:

$$P_D = b_{c_i} - (m_{c_i} * A_{NEW})$$

where b_{c_i} and m_{c_i} are the slope and intercept values associated with the appropriate one of the 2^{n+m} classes of service (step 412). As shown in Fig. 5, the probability of discard changes linearly with changes in the weighted average queue depth. A given packet is discarded when the probability of discard P_D exceeds a "random" number that is produced by a pseudo random generator 25 (steps 414-416). When the weighted average is relatively low, the probability of discard is larger, and thus, the packet is more likely to be discarded.

The slope and intercept values m_{c_i} and b_{c_i} are selected based on trade-offs between keeping links through the router 16 busy and reserving space in the buffer 24 to handle bursts. For higher classes of service the slope and intercept values are selected to be relatively low - such that the probability of discard is low over the entire range from MAX_{c_i} to MIN_{c_i}. The slope and intercept values for the lower classes of service are typically larger, reflecting the greater associated probability of packet loss for the class and the reservation of spaces in the buffer for bursts of packets assigned to the higher classes. The various threshold values, and slope and intercept values are selected such that packet order and probabilities of packet loss are maintained across the 2^{n+m} classes of service.

In prior known routers, implementing a WRED scheme required maintaining average queue depths for all of the classes of service queues at each of the output ports. Thus, for a router to support 8 classes of service over "y" output ports, it had to calculate average queue depths for 8*y separate queues. In the current router 16, the WRED processor calculates the average depth of a single free queue, regardless of the number of classes of service.

A scheduler 26 implements a 2ⁿ class-based weighted round robin (WRR) scheduling scheme for each output port (step 418). The scheduler associates an

appropriate weighting factor W_{Q_j} with each class of service per output port queue. The scheduler de-queues W_{Q_j} packets for transfer from the Q_j queue associated with one of the 2^n classes of service, and then de-queues $W_{Q_{j+1}}$ packets from the Q_{j+1} queue for subsequent transfer. If the Q_{j+1} queue is empty, the scheduler de-queues an appropriate number of packets from the Q_{j+2} queue, and so forth. The scheduler 26 thus ensures that each one of the 2^n classes of service is associated with an appropriate maximum delay limit and through-put allocation.

The class of service mapping, modified WRED scheme and WRR scheme in combination ensure that packets are transferred through the router 16 as if the router supported the 2^{n+m} classes of service. The router 16, however, requires less processing and storage overhead than the prior known routers that support the same number of classes, since the router 16 actually supports 2^n classes of service, and uses a single output buffer to do so.

15 What is claimed is:

CLAIMS

1 1. A router for use in routing packets over a network, the router supporting 2^n classes of
2 service and including:

3 A. a plurality of input ports for receiving packets over the network;
4 B. a plurality of output ports for transferring packets over the network;
5 C. a classifier for assigning packets received by the input ports to 2^{n+m} classes of
6 service and mapping the 2^{n+m} classes of service to the 2^n classes of service that are
7 supported by the router, the classifier assigning to the packet one of 2^m associated levels
8 of priority, wherein each level of priority is associated with a different probability of
9 packet loss;

10 D. means for retaining the packets based on probabilities of discard associated
11 with the 2^{n+m} classes of service; and

12 E. scheduling means for transferring the packets through each of the output ports
13 based on the 2^n classes of service.

1 2. The router of claim 1 further including a multiple storage location buffer for retaining
2 packets to be transferred through the output ports, the buffer linking the storage locations
3 that contain packets in class of service per output port queues and linking available
4 storage locations in a free queue.

1 3. The router of claim 2 wherein the means for retaining the packets further includes:
2 i. means for determining a new weighted average depth for the free queue, and
3 ii. means for determining a probability of discard for a given packet if the new
4 weighted average queue depth falls below a predetermined maximum threshold
5 associated with the class of service to which the packet is assigned.

1 4. The router of claim 3 wherein the discard means discards a given packet if the
2 associated new weighted average depth for the free queue falls below a minimum
3 threshold associated with the class of service to which the packet is assigned.

1 5. The router of claim 4 wherein the discard means calculates the probability of discard
2 as $P_d = c - (m * A_{NEW})$ where c is an intercept and m is a slope that is associated with a line
3 that plots average free queue depth versus probability of discard for the class of service to
4 which the packet is assigned, and A_{NEW} is the new weighted average depth of the free
5 queue.

1 6. The router of claim 5 wherein the discard means calculates the new weighted average
2 depth of the free queue as $A_{NEW} = A_{CURRENT} + w(I - A_{CURRENT})$ where w is a weighting
3 factor, I represents the instantaneous depth of the free queue and $A_{CURRENT}$ is the current
4 weighted average depth of the free queue.

1 7. The router of claim 6 wherein the scheduling means selects packets for transfer based
2 on weighting factors associated with the respective 2^n classes of service.

1 8. A router for use in routing packets over a network, the router supporting 2^n classes of
2 service and including:

- 3 A. a plurality of input ports for receiving packets over the network;
- 4 B. a plurality of output ports for transferring packets over the network;
- 5 C. a multiple storage location buffer for retaining packets to be transferred
6 through the output ports;
- 7 D. means for retaining the packets based on probabilities of discard associated
8 with 2^{n+m} classes of service; and
- 1 E. scheduling means for transferring the packets through each of the output ports
2 based on the 2^n classes of service that the router supports.

1 9. The router of claim 8 further including a classifier for:

- 2 i. assigning packets received by the input ports to 2^{n+m} classes of service,

3 ii. associating the packets with the 2^n classes of service that are supported by the
4 router, and

5 iii. assigning to the packet one of 2^m associated levels of priority, wherein each
6 level of priority is associated with a different probability of packet loss.

1 10. The router of claim 9 wherein the means for retaining the packets further includes

2 i. means for determining a new weighted average queue depth for a free queue
3 that links available buffer storage locations, and

4 ii. means for determining a probability of discard for a given packet if the new
5 weighted average free queue depth falls below a predetermined maximum threshold
6 associated with the class of service to which the packet is assigned.

1 11. The router of claim 10 wherein the discard means calculates the probability of
2 discard as $P_d = c - (m * A_{NEW})$ where c is an intercept and m is a slope that are associated
3 with a line that plots average free queue depth versus probability of discard for the class
4 of service to which the packet is assigned, and A_{NEW} is the new weighted average depth
5 of the free queue.

1 12. The router of claim 11 wherein the discard means calculates the new depth of the free
2 queue as $A_{NEW} = A_{CURRENT} + w(I - A_{CURRENT})$ where w is a weighting factor, I represents
3 the instantaneous depth of the free queue and $A_{CURRENT}$ is the current weighted average
4 depth of the free queue.

1 13. The router of claim 12 wherein the discard means discards a given packet if the new
2 weighted average free queue depth falls below a minimum threshold associated with the
3 class of service to which the packet is assigned.

1 14. The router of claim 12 wherein the discard means retains a given packet if the new
2 weighted average free queue depth is above a maximum threshold associated with the
3 class of service to which the packet is assigned.

1 15. The router of claim 8 wherein the scheduling means selects packets for transfer
2 through each output port based on weighting factors associated with the respective 2^n
3 classes of service.

1 16. The router of claim 15 wherein the buffer links retained packets in class of service
2 per output port queues and the scheduling means selects packets from the class of service
3 per output port queues.

1 17. A method of routing packets through a router that supports 2^n classes of service, the
2 method including the steps of:
3 A. receiving packets through one or more input ports;
4 B. assigning packets received by the input ports to 2^{n+m} classes of service and
5 mapping the 2^{n+m} classes of service to the 2^n classes of service that are supported by the
6 router, the classifier assigning to the packet one of 2^m associated levels of priority,
7 wherein each level of priority is associated with a different probability of packet loss;
8 C. retaining the packets based on probabilities of discard associated with the 2^{n+m}
9 classes of service; and
10 D. transferring the packets through one or more output ports based on the 2^n
11 classes of service.

1 18. The method of routing packets of claim 17 further including in the step of retaining
2 the packets the steps of :
3 i. retaining the packets in a multiple storage location buffer and linking available
4 storage locations to a free queue,
1 ii. determining a new weighted average depth for the free queue, and
2 iii. determining a probability of discard for a given packet if the new weighted
3 average queue depth falls below a predetermined maximum threshold associated with the
4 class of service to which the packet is assigned.

- 1 19. The method of routing packets of claim 18 including in the step of retaining the
- 2 packets the further step of discarding a given packet if the new weighted average depth
- 3 for the free queue falls below a minimum threshold associated with the class of service to
- 4 which the packet is assigned.

- 1 20. The method of routing packets of claim 19 wherein the step of retaining the packets
- 2 includes calculating the probability of discard as $P_d = c - (m * A_{NEW})$ where c is an intercept
- 3 and m is a slope associated with a line that plots weighted average free queue depth
- 4 versus probability of discard for the class of service to which the packet is assigned, and
- 5 A_{NEW} is the new weighted average depth of the free queue.

- 1 21. The method of routing packets of claim 20 wherein the step of retaining the packets
- 2 includes calculating the new weighted average depth of the free queue as $A_{NEW} =$
- 3 $A_{CURRENT} + w(I - A_{CURRENT})$ where w is a weighting factor I represents the instantaneous
- 4 depth of the free queue and $A_{CURRENT}$ is the current weighted average queue depth.

- 1 22. The method of claim 21 wherein the discard means retains a given packet of the new
- 2 weighted average free queue depth is above a maximum threshold associated with the
- 3 class of service to which the packet is assigned.

- 1 23. The method of routing packets of claim 17 wherein the step of transferring packets
- 2 through the more or more output port transfers the packets based on weighting factors
- 3 associated with the respective 2^n classes of service.

- 1 24. A method of routing packets through a router that supports 2^n classes of service, the
- 2 method including:
 - 3 A. receiving packets through one or more input ports and assigning the packets to
 - 4 2^{n+m} classes of service;

5 B. retaining packets based on probabilities of discard associated with the 2^{n+m}
6 classes of service in a multiple storage location buffer that links available storage
7 locations to a free queue; and

1 C. transferring the packets through one or more output ports based on the 2^n
2 classes of service.

1 25. The method of routing of claim 24 further including the steps:

2 i. associating the packets that are assigned to the 2^{n+m} classes of service with the
3 2^n classes of service that are supported by the router, and
4 ii. assigning to the respective packets one of 2^m associated levels of priority,
5 wherein each level of priority is associated with a different probability of packet loss.

1 26. The method of routing packets of claim 25 wherein the step of retaining the packets
2 includes:

3 a. determining a new weighted average depth for the free queue, and
4 b. determining a probability of discard for a given packet if the new weighted
5 average free queue depth falls below a predetermined maximum threshold associated with
6 the class of service to which the packet is assigned.

1 27. The method of routing packets of claim 26 wherein the step of retaining packets
2 further includes calculating the probability of discard as $P_d = c - (m * A_{NEW})$ where c is an
3 intercept and m is a slope that are associated with a line that plots average free queue
4 depth versus probability of discard for the class of service to which the packet is assigned,
5 and A_{NEW} is the new weighted average depth of the free queue.

1 28. The method of routing packets of claim 27 wherein the step of retaining packets
2 further includes calculating the new weighted average depth of the free queue as $A_{NEW} =$
3 $A_{CURRENT} + w(I - A_{CURRENT})$ where w is a weighting factor, I represents the instantaneous
4 depth of the free queue and $A_{CURRENT}$ is the current weighted average queue depth.

1 29. The method of routing packets of claim 26 wherein the step of retaining packets
2 further includes discarding a given packet if the new weighted average free queue depth
3 falls below a minimum threshold associated with the class of service to which the packet
4 is assigned.

1 29. The method of routing packets of claim 26 wherein the step of retaining packets
2 further includes retaining a given packet if the new weighted average free queue depth is
3 above the maximum threshold associated with the class of service to which the packet is
4 assigned.¹

ABSTRACT OF THE DISCLOSURE

A router includes a classifier that classifies packets assigned to 2^{n+m} classes of service into 2^n classes of service that are supported by the router. The classifier then sets the loss priorities of the respective packets to one of m levels. The router uses a modified weighted random early detection scheme that is based on probabilities of discard associated with the 2^{n+m} classes of service to determine whether to retain or discard particular packets. The router uses a single buffer to store packets directed to all of the various output ports. The available storage locations in the buffer are linked to a free queue and a weighted average depth of the free queue, is used to determined whether or not to retain a given packet. The router compares the weighted average depth of the free queue to maximum and minimum thresholds associated with the particular 2^{n+m} class of service to which the packet is assigned. If the weighted average is above the maximum threshold, the packet is retained. If the weighted average is below the minimum threshold, the packet is discarded. If the weighted average is between the two thresholds, a probability of discard that is based on the 2^{n+m} classes of service is calculated and compared to a random value to determine whether or not the packet should be retained. If the probability of discard exceeds the random value, the packet is discarded. Otherwise, the packet is retained. The retained packets are then transferred through the various output ports of the router based on weighting factors associated with the 2^n classes of service.




FIG. 1





FIG. 2

CLASSES OF SERVICE	COS TAGS	CLASS OF SERVICE PER OUTPUT PORT QUEUES
IS/ISLL GUARANTEED	000/001	Q_0
ROVING MANAGEMENT	000/001	Q_0
IS/ISLL CONTROLLED LOAD	010/011	Q_1
RESERVED	100/101	Q_2
BEST EFFORT	110/111	Q_3

FIG. 3

FIG. 4






FIG. 5

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below-named inventor, I hereby declare that:

My residence, post-office address, and citizenship are as stated below next to my name.

I believe I am an original, first, and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled ROUTER WITH CLASS OF SERVICE MAPPING, the specification of which was filed on December 31, 1997, and accorded Serial Number 09/900,478 .

I hereby state that I have reviewed and understand the contents of the above-identified application specification, including the claims, as amended by any amendment specifically referred to herein.

I acknowledge the duty to disclose all information known to me that is material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code §119(a)-(d) of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate filed by me on the same subject matter having a filing date before that of the application on which priority is claimed: None.

I hereby claim the benefit under Title 35, United States Code §119(e) of the following U.S. provisional application: None.

I hereby claim the benefit under Title 35, United States Code §120, of the United States Application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United State Code, §112, I acknowledge the duty to disclose all information that is material to patentability in accordance with Title 37, Code of Federal Regulations, §1.56, and which became available to me between the filing date of the prior application and the national or PCT international filing date of this application: None.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment or both under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint Michael E. Attaya, Reg. No. 31,731; Charles J. Barbas, Reg. No. 32,959; Joseph H. Born, Reg. No. 28,283; Robert A. Cesari, Reg. No. 18,381; Steven J. Frank, Reg. No. 33,497; Christopher K. Gagne, Reg. No. 36,142; William A. Loginov, Reg. No. 34,863; Robert A. Mazzarese, Reg. No. P42,852; John F. McKenna, Reg. No. 20,912; Martin J. O'Donnell, Reg. No. 24,204; Thomas C. O'Konski, Reg. No. 26,320; Michael R. Reinemann, Reg. No. 38,280; Rita M. Rooney, Reg. No. 30,585; Heather B. Shapiro, Reg. No. 41,305; and Patricia A. Sheehan, Reg. No. 32,301, Cesari and McKenna, LLP, 30 Rowes Wharf, Boston, Mass. 02110, jointly, and each of them severally, my attorneys and attorney, with full power of substitution, delegation and revocation, to prosecute this application, to make alterations and amendments therein, to receive the patent and to transact all business in the Patent and Trademark Office connected therewith. Please direct all telephone calls to Patricia A. Sheehan at (617) 951-2500. Please address all correspondence to Patricia A. Sheehan.

Arthur Lin 5-20-98
Arthur Lin Date

Residence: 211 Firethorn Court
San Ramon, CA 94583

Citizenship United States

Post Office Address: Same as above

Thomas J. Edsall 5-13-98
Thomas J. Edsall Date

Residence: 471 Mountain View Avenue
Mountain View, CA 94041

Citizenship United States of America

Post Office Address: Same as above

Tauseef Hashmi
Tauseef Hashmi

May 13, 1998
Date

Residence: 44263 Lupine Place
Freemont, CA 94539

Citizenship United States

Post Office Address: Same as above

H:\112025\0067\PROSECUT\DECLARE.DOC